

DFG-SNF Research Group FOR916

Statistical Regularization and Qualitative Constraints

Thomas Hotz
Laurie Davies

Philipp Marnitz
Zakhar Kabluchko

Rahel Stichtenoth
Axel Munk

Locally adaptive image denoising by a statistical multiresolution criterion

Preprint FOR916 09-08

Preprint-Series of the Research Group FOR916

Locally adaptive image denoising by a statistical multiresolution criterion

Thomas Hotz, Philipp Marnitz, Rahel Stichtenoth,
Laurie Davies, Zakhar Kabluchko and Axel Munk

November 2009

Abstract

We demonstrate how one can choose the smoothing parameter in image denoising by a statistical multiresolution criterion, both globally and locally. Using inhomogeneous diffusion and total variation regularization as examples for localized regularization schemes, we present an efficient method for locally adaptive image denoising. As expected, the smoothing parameter serves as an edge detector in this framework. Numerical examples illustrate the usefulness of our approach. We also present an application in confocal microscopy.

Keywords: Image reconstruction, statistical multiresolution criterion, bandwidth selection.

2000 Mathematics Subject Classification: Primary 68U10, 62G08; Secondary 60G70.

1 Introduction

Image denoising is one of the main tasks in image analysis, as documented by numerous articles and books published on the subject, cf. e.g. (Scherzer et al. 2009), (Buades et al. 2005) and (Aubert and Kornprobst 2002). Accordingly, statisticians have contributed their share: using probabilistic models for the true image, Bayesian methods were among the first to add a statistical perspective to the subject, see e.g. (Geman and Geman 1984), (Besag 1986), and (Winkler 2003). Taking a frequentist's point of view, image denoising

becomes a smoothing or reconstruction problem, see e.g. (Hall and Titterton 1986) and (Polzehl and Spokoiny 2000, 2003) as well as (Korostelev and Tsybakov 1993). Acknowledging the non-smooth nature of many images caused by sharp edges, cf. (Chu et al. 1998) and (Donoho 1999), prompted a generalization of the well-established smoothing techniques, e.g. drawing on methods from one-dimensional jump detection, see (Qiu 2005, 2007); cf. also (Mrázek et al. 2006) for a unifying framework for many popular numerical and statistical denoising techniques.

Put simply, *statistical image denoising* amounts to reconstructing a noiseless image f given a noisy image y . Usually, one assumes the noise ϵ to be additive, i.e. $y = f + \epsilon$. In this article, we assume that the noise is generated at random; more specifically, we assume for pixels (i, j) that

$$y_{ij} = f_{ij} + \epsilon_{ij} \tag{1}$$

with ϵ_{ij} identically and independently distributed Gaussian random variables with zero mean and variance σ^2 , i.e.

$$\epsilon_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2); \tag{2}$$

in particular we assume the value at each pixel to be a real number, i.e. $y_{ij}, f_{ij}, \epsilon_{ij} \in \mathbf{R}$. This models a grey-scale image, though in practice its grey levels are usually restricted to a finite number of discrete values, e.g. to integers between 0 and 255. In many applications, a Gaussian assumption on the noise is therefore not very plausible but other noise processes will be more suitable. Nonetheless, for simplicity's sake we will lay out our basic ideas under a Gaussian assumption; possible extensions beyond this are briefly discussed at the end of Section 3. We note, however, that for a well-calibrated image y making good use of the range of (not too few) possible values, this assumption is not very crucial, as long as the errors are i.i.d., symmetric and feature no heavy tails: the true image f will then take values well inside the range, itself not necessarily being restricted to discrete values – a property it passes onto the errors ϵ_{ij} .

Figure 1 shows an artificial example where f exhibits varying degrees of smoothness (left), and Gaussian white noise has been added to obtain the data y (right). The assumption about the noise can then be exploited in order to distinguish between the ‘true’, noiseless image f and the noise ϵ as demonstrated in the following sections.




Figure 1: (a) 256×256 pixel test image f taking values in $[0, 5]$; dashed line indicates row 64 shown in detail in Figure 8. (b) Simulated data y with noise level $\sigma = 1$.

Clearly, this is impossible unless assumptions are made about f , e.g. that it varies slowly from pixel to pixel. In order to be able to formulate such ‘smoothness’ assumptions more precisely, we let (in a slight abuse of notation) f_{ij} denote the values of a function f on a regular grid, i.e. $f : [0, 1]^2 \rightarrow \mathbf{R}$ with $f_{ij} = f(\frac{i}{n}, \frac{j}{n})$ where $i, j \in \{1, \dots, n\}$. Note that we use a square grid for ease of notation, with $x_{ij} = (\frac{i}{n}, \frac{j}{n})$ in the following; one easily can extend all the analyses to rectangular grids, to higher dimensions, or to non-uniform sampling schemes through the use of finite elements, if required, cf. e.g. (Ern and Guermond 2004). Now that f can be viewed as a function, ‘smoothness’ can be defined more rigorously to mean that f belongs to some function class, e.g. that f is in a Sobolev or Besov ball, or that f has bounded variation, cf. e.g. (Korostelev and Tsybakov 1993).

Image denoising techniques like the ones we discuss in Section 2 generally require the choice of a *smoothing* or *regularization parameter* a which determines how much smoothing is to be applied. This parameter might be *localized*, allowing for different amounts of smoothing to be applied to different parts of the image; the regularization parameter then becomes a function $a : [0, 1]^2 \rightarrow \mathbf{R}$. Since one might want to smooth more where the true image f is smoother, this then enables one to adapt to differing levels of smoothness across the image. The reconstruction or denoised image \hat{f} hence depends on that smoothness parameter. The aim of this research was to find

a purely data-driven and generally applicable way to choose a . This will be illustrated with two specific denoising techniques, namely for linear diffusion as well as for TV penalization. We stress, however, that our approach is in principle applicable to any regularization technique which depends on properly choosing a regularization parameter, the latter possibly being a function as described above.

The main idea can be summarized as follows: consider the residuals $r_{ij} = y_{ij} - \hat{f}_{ij}$; they depend on the smoothness parameter a . Indeed, if we smooth too much some structures which were present in f have been smoothed away – these structures then will be left in the residuals. Had we found the perfect reconstruction, i.e. for $\hat{f} = f$, however, the residuals form white noise, see (1). One possible way to decide whether we smoothed too much is therefore to check whether the residuals look like white noise – if there is still some structure left in them we must have smoothed too much. We note that this idea is at the heart of statistical methods for automatically selecting the regularization parameter.

As the key ingredient for choosing the smoothness parameter a we propose a *statistical multiresolution criterion*. Introduced in Section 3, it measures deviations from the hypothesis that the residuals are white noise. An important feature of this criterion is that it not only detects if the residuals deviate from white noise but also *where*. This then allows for a locally adaptive choice of a . In Section 4 we derive an algorithm for a data-driven selection of a . Numerical details and results are given in Section 5, alongside an example from confocal microscopy. Finally, we summarize and discuss what we have achieved in Section 6. The interested reader might also find the Ph.D. thesis of Stichtenoth (2007) useful, on which parts of this article have been based.

While there exists an extensive literature on localized, data-driven ways of choosing the smoothing parameter in one-dimensional function estimation, see e.g. (Lepskii et al. 1997), (Gijbels and Mammen 1998) and (Dümbgen and Spokoiny 2001), and several articles have been published on multiresolution criteria for one-dimensional settings, cf. the references in Section 3, higher-dimensional situations have scarcely been treated. In fact, to the authors' knowledge (Bissantz et al. 2006, 2008) and (Davies and Meise 2008) are the only published work on a data-driven choice of the smoothing parameter by statistical multiresolution techniques in higher dimensions; however, Bissantz et al. (2006, 2008) essentially still apply a one-dimensional version of the multiresolution criterion to this end whereas the two-dimensional ap-

plication in (Davies and Meise 2008) is described very briefly and somewhat rudimentary.

The present paper therefore appears to be the first to give a comprehensive and detailed exposition of a two-dimensional multiresolution criterion, and also in proposing a method for choosing a localized smoothing parameter in a purely data-driven and general fashion. We claim that it can be applied to a wide range of reconstruction techniques depending on a global or local regularization parameter; this will be demonstrated exemplarily for linear diffusion filtering and total variation regularization below.

Let us now consider these commonly used methods for image reconstruction, i.e. for noise removal.

2 Image denoising

Among the best studied denoising technique in image processing is the *linear diffusion filter*, see e.g. (Weickert 1998) or (Guichard et al. 2004), which has the physical interpretation of an evolution of the *heat equation*. For constant diffusivity $a \in \mathbf{R}$, it is given by

$$\begin{cases} \frac{\partial}{\partial t} u_a(x, t) = a \Delta_x u_a(x, t) \\ u_a(x, 0) = y \end{cases} \quad (3)$$

which is solved by the convolution of the initial data $y : \mathbf{R}^2 \rightarrow \mathbf{R}$ with a Gaussian kernel, i.e.

$$u_a(x_*, t) = \int K_{\sqrt{2at}}(x - x_*) y(x) dx \quad (4)$$

with

$$K_h(x) = \frac{1}{2\pi h} \exp\left(-\frac{\|x\|^2}{2h^2}\right); \quad (5)$$

it is therefore also called *Gaussian filtering*. This diffusion filter is very well understood by now: it is a low-pass filter effectively reducing white noise; the heat equation is the limit of repeated convolving with any kernel fulfilling some moment conditions (Guichard et al. 2004, Ch. 2); it gives rise to a linear scale space with many desirable properties (Chaudhuri and Marron 2000); and there are many sets of such properties which uniquely determine this scale space, see e.g. (Witkin 1983), (Koenderink 1984), (Alvarez et al. 1993), or (Weickert 1998, Table 1.1) for an overview.

Note that the aim of this paper is not to propagate this specific filter but rather to illustrate with this particular example how statistical multiresolution analysis can be used to effectively choose the diffusivity, or some other regularization parameter, globally or locally – even in the case when the optimal choice depends on the smoothness of the true image f ; other regularization techniques will be discussed at the end of this section.

In statistics, model (1), where f is not assumed to have a certain low-dimensional, parametric form but instead to belong to some non-parametric function class, is called a (two-dimensional) *nonparametric regression model*. A standard approach to estimate f from the data y is to use a *kernel estimator* \hat{f}_K : choosing some kernel $K : \mathbf{R}^2 \rightarrow \mathbf{R}$, one forms a local average where the kernel defines the weights,

$$\hat{f}_K(x_*) = \frac{\sum_{ij} K(x_{ij} - x_*) y_{ij}}{\sum_{ij} K(x_{ij} - x_*)}, \quad (6)$$

for arbitrary $x_* \in [0, 1]^2$. A popular choice for K is the isotropic Gaussian kernel K_h with bandwidth h given in (5). We refer to (Wand and Jones 1995) for a broader treatment of kernel estimation, mentioning just one of the many textbooks on the subject.

From (4) we see that, after discretization in space and ignoring boundary effects, the solution of the heat equation at time t is the kernel estimator with bandwidth $h = \sqrt{2at}$, i.e. the rôle the (global) bandwidth is playing in kernel estimation is played by time when smoothing with the heat equation. Alternatively, we can vary the diffusivity and fix the endpoint in time since $u_1(x, t) = u_t(x, 1)$. We will take the latter approach, calling the process *homogeneous diffusion*, with the diffusivity $a \in \mathbf{R}$ to be specified. Finally, we also discretize time to only two time points, namely 0 and 1, thereby obtaining a one-step approximation to the heat equation which is computed on the same grid as the data are on. For $a \in \mathbf{R}$, this defines an estimator $\hat{f}_{\text{hom.diff}}$ through

$$\hat{f}_{\text{hom.diff}} - y = a \tilde{\Delta}_x \hat{f}_{\text{hom.diff}} \quad (7)$$

where $\tilde{\Delta}_x$ is a discretization of the Laplacian in space.

Clearly, the choice of the diffusivity – or, equivalently, of the bandwidth – is critical for the performance of the estimator: a large diffusivity will reduce the variance of the estimator but at the expense of increasing its bias. While the variance depends on the noise level σ , the smoothness of the function influences the bias: the smoother the function, the less the bias, with




Figure 2: Homogeneous diffusion applied to the data in Figure 1(b), with diffusivity chosen as $a = 40$ in (a), and $a = 0.1$ in (b).

constant functions giving no bias at all. This is illustrated in Figure 2: while a large diffusivity (left) is permissive in regions where the underlying signal, see Figure 1(a), is smooth, e.g. in the interior of the large circle or within the background, sharp edges and smaller features like the circles in the lower right or the “valleys” in the lower left are smoothed away; the latter need a lower diffusivity (right) but this then results in unnecessary and unwanted undersmoothing in the upper left part. The optimal diffusivity strikes a balance between bias and variance – a goal which is difficult to achieve, given that the bias depends on the smoothness of the *unknown* function. An estimator that chooses the diffusivity in a purely data-driven manner will be called *adaptive*.

Moreover, instead of choosing a global, one-fits-all, diffusivity a , one would rather want to choose it locally, as the function’s smoothness is a local feature, too, cf. again Figure 2. Estimators that choose the diffusivity locally in dependence on the data observed are *locally adaptive*. We stress that in the context of our work “adaptivity” is not meant in the sense that the estimator attains certain optimality properties over scales of spaces (see e.g. (Lepskii 1990), (Tsybakov 1998) and (Dümbgen and Spokoiny 2001)), but we rather use this term heuristically to express that the estimator chooses e.g. its diffusivity in a data-driven fashion, possibly taking local smoothers of f into account. Deriving such an estimator is the aim of this paper. It will be achieved by the use of the *statistical multiresolution criterion* which allows to

statistically decide whether some reconstruction's residuals still contain parts of the signal. We defer to Section 3 for a longer discussion of its analytical properties and to Section 4 for its application to image reconstruction.

The task of choosing the bandwidth locally is similar to selecting a spatially varying diffusivity $a(x)$ in (3), and correspondingly in (7). We point out that, nonetheless, these techniques are not equivalent; this can easily be seen if one compares the effect of a zero diffusivity along a curve which inhibits the exchange of information across even if the diffusivity is large close to that curve. Hence, local diffusivities allow to respect sharp edges. Note that in the case of a local diffusivity $a : [0, 1]^2 \rightarrow \mathbf{R}$, equation (7) can still be solved very efficiently as opposed to the computation of kernel estimators with varying bandwidths, see Section 5 for details. We thus obtain an estimator $\hat{f}_{\text{inhom.diff.}}$ through this *inhomogeneous diffusion* process. The corresponding linear operator will be denoted by L_a , i.e.

$$\hat{f}_{\text{inhom.diff.}} = L_a y = a \tilde{\Delta}_x L_a y + y. \quad (8)$$

For details on how to solve this equation as well as on differences to local bandwidth selection, we refer to (Stichtenoth 2007).

The variational formulation of the continuous version of (8) is given by

$$\operatorname{argmin}_{g \in H^2} \left\| \frac{g - y}{a} \right\|^2 + \int |\nabla g|^2 \quad (9)$$

where $\|\cdot\|$ and $|\cdot|$ denote the L_2 -norm and the Euclidean norm in \mathbf{R}^2 , respectively. Here, H^2 denotes the Sobolev space of functions in L_2 having weak second derivatives also in L_2 , i.e. $H^2 = \{g : \sum_{|\alpha| \leq 2} \|D^\alpha g\| < \infty\}$. Clearly, the first term in (9) is a weighted data fit whereas the second term is a smoothness penalty; this allows to view a also as local weights, or being proportional to a locally varying standard deviation σ of the errors.

Other penalties can also be thought of: considering the poor performance of the kernel estimator at sharp edges, one might also be interested in a total variation (TV) penalty as introduced by Rudin et al. (1992),

$$\hat{f}_{\text{TV}} = \operatorname{argmin}_{g \in BV} \left\| \frac{g - y}{a} \right\|^2 + \int |\nabla g|, \quad (10)$$

where the second term symbolically stands for the TV semi-norm, i.e. $\int |\nabla g| := \sup_{\phi \in C_c^1(\mathbf{R}^n), \|\phi\|_\infty \leq 1} \int g \operatorname{div} \phi$ where $\phi \in C_c^1(\mathbf{R}^n)$ iff $\phi \in C^1(\mathbf{R}^n)$ and ϕ has

compact support. We also assume a discretization in space as above; as usual, BV denotes the space of functions with bounded variation, defined as $BV = \{g \in L_1 : \int |\nabla g| < \infty\}$.

The main question that needs to be answered is how to choose the diffusivity (or weights) a . This will be addressed in the following two sections. We note that it is not always possible to localize the regularization parameter. One class of denoising techniques for which this is the case are iterative methods; there, the iteration at which to stop the algorithm plays the rôle of the regularization parameter, see e.g. (Bissantz et al. 2006, 2008) for the EM algorithm, a.k.a. Richardson-Lucy algorithm, or (Bissantz et al. 2007) for a general treatment of iterative regularization schemes. We therefore also discuss how to select a globally, although the focus will be on a local selection strategy.

3 A statistical multiresolution criterion

We have assumed the errors ϵ_{ij} to be white noise, see (2). Clearly, considering the residuals $r_{ij} = y_{ij} - \hat{f}_{ij}$ for some estimator \hat{f} , we would want them to be as close to white noise as possible – if there is any structure left in the residuals then the estimator must have missed essential features of the true f . In this section, we therefore aim to derive a statistical test for the hypothesis that the residuals are white noise, i.e. $r_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$, against the alternative that there is some structure left in them. To this end we define random variables

$$\omega_P = \frac{1}{\sqrt{\#\{x_{ij} \in P\}}} \sum_{x_{ij} \in P} r_{ij} \quad (11)$$

for suitable subsets $P \subseteq [0, 1]^2$ of the image domain, to be specified later. We call ω_P the *multiresolution coefficient* of P . The collection \mathcal{P}_n of subsets P cannot be too large and will typically be of polynomial order $O(n^k)$ of the number of observations n . This is the case if \mathcal{P}_n is a Vapnik-Cervonenkis (VC) class of subsets, cf. (Pollard 1984; Devroye and Lugosi 2001). In fact the family \mathcal{P}_n must be a VC-class for the procedure to make sense as otherwise too many subsets are picked out and the stochastic fluctuation of the random variables ω_P can no longer be controlled simultaneously over all subsets $P \in \mathcal{P}_n$. The choice of \mathcal{P}_n is subtle as it influences the limit behaviour; it will be addressed later where we propose amongst others a dyadic squares partitioning.

Note that $\omega_P \sim \mathcal{N}(0, \sigma^2)$ under the hypothesis. If however there is some information left in the residuals then the mean of some residuals will no longer be 0. Hence the corresponding multiresolution coefficients ω_P are also no longer distributed around 0 but around the mean of the signal left in the residuals, thus becoming large in absolute value in comparison to the expected behaviour of white noise.

Given an appropriate collection of subsets \mathcal{P}_n of $[0, 1]^2$ and $\alpha, 0 < \alpha < 1$, we can determine (e.g. through simulations) some $\tau_n = \tau_n(\alpha)$ such that

$$\mathbf{P} \left(\max_{P \in \mathcal{P}_n} \frac{|\sum_{x_{ij} \in P} Z_{ij}|}{\sqrt{\#\{x_{ij} \in P\}}} \leq \sqrt{\tau_n \log n^2} \right) = \alpha, \quad (12)$$

where Z_{ij} are standard Gaussian white noise random variables. For any function $g : [0, 1]^2 \rightarrow \mathbf{R}$ and $P \in \mathcal{P}_n$ we define

$$w(g, P) = \frac{\sum_{x_{ij} \in P} (y_{ij} - g_{ij})}{\sqrt{\#\{x_{ij} \in P\}}} \quad (13)$$

where $g_{ij} = g(i/n, j/n)$ and put

$$\mathcal{F}_n = \mathcal{F}_n(\mathcal{P}_n, \sigma, \tau_n) = \left\{ g : \max_{P \in \mathcal{P}_n} |w(g, P)| \leq \sigma \sqrt{\tau_n \log n^2} \right\}. \quad (14)$$

For data generated under (1) with $\epsilon_{ij} = \sigma Z_{ij}$ it is seen that \mathcal{F}_n is an exact, universal and non-asymptotic confidence region for any $f : [0, 1]^2 \rightarrow \mathbf{R}$:

$$\mathbf{P}(f \in \mathcal{F}_n) = \alpha$$

provided this event is measurable, see (Davies et al. 2009). The confidence region \mathcal{F}_n contains many functions which are of little interest, for example, all functions which interpolate the data. This may be seen as a severe case of overfitting. In general, however, we are interested in simple or, if possible, the simplest functions in \mathcal{F}_n where simplicity may be defined in terms of smoothness, shape or sparsity or some combination of all three. Regularization within \mathcal{F}_n leads to optimization problems such as

$$\text{minimize } TV(g^{(k)}) \quad \text{subject to } g \in \mathcal{F}_n \quad \text{for some } k = 0, 1, \dots \quad (15)$$

where $TV(g^{(k)})$ denotes some definition of total variation of the function $g^{(k)}$. For examples of this approach for one-dimensional data we refer to (Mammen and van de Geer 1997) and (Davies et al. 2009).

In some cases where the optimization problem is algorithmically too difficult to be solved we may nevertheless have a sequence $\tilde{f}_m, m = 1, 2, \dots$ of good candidate functions of increasing complexity. This is the case we consider in this paper and the strategy consists of choosing the first function \tilde{f}_m which lies in \mathcal{F}_n . The success of this strategy depends largely on how good the candidate functions are.

The definition of \mathcal{F}_n involves σ which has to be estimated from the data. We propose a robust estimator for this purpose in Section 4, see (28). We note that it would be possible to refine the simple substitution of $\hat{\sigma}_n$ for σ slightly so that \mathcal{F}_n becomes an honest (Li 1989; Genovese and Wasserman 2008), universal and non-asymptotic confidence region for f , i.e.

$$\mathbf{P}(f \in \mathcal{F}_n) \geq \alpha.$$

As already mentioned $\tau_n(\alpha)$ can be obtained by simulations. However, if the asymptotic behaviour of τ_n can be determined then it is often sufficient to use $\tau_\infty = \lim_{n \rightarrow \infty} \tau_n$. This is not a simple problem and we consider it in more detail below. In the cases we consider we have $\tau_\infty = 2$.

Following these considerations, we define the *statistical multiresolution criterion* M by

$$M_n = \frac{1}{\sqrt{2 \log n^2}} \max_{P \in \mathcal{P}_n} |\omega_P| = \frac{1}{\sqrt{2 \log n^2}} \max_{P \in \mathcal{P}_n} \frac{|\sum_{x_{ij} \in P} r_{ij}|}{\sqrt{\#\{x_{ij} \in P\}}}, \quad (16)$$

for some collection \mathcal{P}_n of subsets P , the normalization factor $\sqrt{2 \log n^2}$, or $\sqrt{2 \log n^d}$ in a d -dimensional setting, becoming clear in a short while. Note that the choice of \mathcal{P}_n is subtle as it influences the limit behaviour; it will be addressed later where we propose to use a dyadic squares partitioning.

The multiresolution criterion in the form of (16) was introduced by Siegmund and Venkatraman (1995) to detect change points; Davies and Kovac (2001) were first to use it for one-dimensional non-parametric regression, Bisantz et al. (2006, 2008) applied it to positron emission tomography, while a similar criterion has been introduced by Dümbgen and Spokoiny (2001) as well as Dümbgen and Walther (2008) in the context of testing qualitative hypotheses in non-parametric regression. A major difference between (16) and the latter authors' multiscale statistic is that they calibrate the average residuals by a term of the order $\sqrt{\log(n/\#P)}$ in order to enhance medium and large scales. In many imaging problems, however, the features on the

small scales are most important as they reflect rapid local change at edges. This particularly results in a completely different limiting behaviour, compare Theorem 2 in (Dümbgen and Walther 2008) with Theorem 1 below. Siegmund and Worsley (1995) test for the presence of a signal of unknown scale and position in arbitrary dimensions, using Gaussian weights in their multiresolution criterion; they derive its asymptotic distribution and discuss its power. Davies and Meise (2008) used this criterion in one- and two-dimensional settings to determine the weights of smoothing splines.

Although the ω_P are identically distributed they are dependent, rendering the distribution of M_n difficult to obtain analytically. Recently, Kabluchko and Munk (2008) established a.s. convergence, i.e.

$$M_n \rightarrow \sigma \text{ a.s. for } n \rightarrow \infty \quad (17)$$

for \mathcal{P}_n the collection of all squares and rectangles; for the one-dimensional case and intervals this was shown by Shao (1995). For the latter case, Siegmund and Venkatraman (1995) proved that M_n is asymptotically Gumbel-distributed, cf. also (Kabluchko 2007); (Siegmund and Yakir 2000) suggests this also to hold in the multi-dimensional case. We will present a similar result for the collection of dyadic cubes in Theorem 2 on page 15. It will be based on the following, general theorem which is proved in the appendix:

Theorem 1. *For each $N \in \mathbb{N}$, let $(\xi_1^{(N)}, \dots, \xi_N^{(N)})$ be a Gaussian vector with standardized marginal distributions. Suppose that for every $\varepsilon > 0$ and some constant $\rho < 1$ not depending on N we have*

$$\#\{(i, j) \in \{1, \dots, N\}^2 : \text{Cov}(\xi_i^{(N)}, \xi_j^{(N)}) \neq 0\} = O(N^{1+\varepsilon}) \text{ as } N \rightarrow \infty \quad (18)$$

and

$$|\text{Cov}(\xi_i^{(N)}, \xi_j^{(N)})| \leq \rho \text{ provided that } i \neq j. \quad (19)$$

Then

$$\lim_{N \rightarrow \infty} \mathbf{P} \left[\max_{i=1, \dots, N} \xi_i^{(N)} \leq a_N + b_N \tau \right] = \exp(-e^{-\tau}) \quad (20)$$

where a_N and b_N are sequences of constants defined by

$$a_N = \sqrt{2 \log N} + \frac{-1/2 \log \log N - \log 2\sqrt{\pi}}{\sqrt{2 \log N}}, \quad b_N = \frac{1}{\sqrt{2 \log N}}. \quad (21)$$

We note that condition (18) states that the covariance matrix of $(\xi_i^{(N)})_{i=1,\dots,N}$ is “sparse”, that is, it contains at most $O(N^{1+\varepsilon})$ non-zero elements. Condition (19) states that the covariance matrix has mutual coherence less than 1, i.e. its off-diagonal elements are bounded away from ± 1 .

While the exact distribution of M_n might not be available for more general collections \mathcal{P}_n or in a finite setting, a critical value for testing the hypothesis can in principle be obtained through simulation.

We note that extensions to more general error distributions are possible, see (Shao 1995), (Kablichko and Munk 2008) or (Nardi et al. 2008). For the particular case of Poisson data with not too small intensities, we suggest to approximate these by Gaussian distributions: hypothesising that the data y stem from intensities \hat{f} , compute residuals $r_{ij} = \hat{f}_{ij}^{-1/2}(y_{ij} - \hat{f}_{ij})$ which are approximately i.i.d. like Gaussian white noise with standard deviation 1, see also Section 5.

4 Data-driven choice of the diffusivity

We now return to the question of how to choose the diffusivities for the estimators defined in Section 2. In view of the test criterion in (16), we require that the conditions

$$|\omega_P| \leq \sigma \sqrt{\delta \log n^2} \quad (22)$$

hold for all $P \in \mathcal{P}_n$. The asymptotics in (17) suggest $\delta > 0$ to be chosen close to 2, in such a way that the error of the first kind is below a certain, prespecified significance level α , say $\alpha = 5\%$. If all these conditions are fulfilled, we cannot reject the hypothesis that the residuals are white noise, or equivalently that the reconstruction \hat{f} is the true function f ; thence we are to accept \hat{f} as possibly being the true f .

An important property of the inequalities (22) to note is that they are only one-sided, i.e. they are only violated if the multiresolution coefficients ω_P get large in absolute value. This will happen if we oversmooth, using too large a diffusivity. If we smooth too little, using too small a diffusivity, however, then the residuals get too small and so do the ω_P , and the criterion will not be violated. For example, estimating f by $\hat{f} = y$ leads to all residuals r_{ij} being 0, and hence the multiresolution criterion is trivially fulfilled.

Nonetheless, in many situations the collection \mathcal{P}_n together with the candidate sets \mathcal{F}_n constitute a nested sequence of increasing complexity when

τ_n increases in (14); then it is reasonable to choose $\hat{f}_n \in \mathcal{F}_n$ such that the maximum of the left hand side of (22) is as close as possible to equality. Bissantz et al. (2006, 2008) have investigated this strategy for stopping the EM algorithm in positron emission tomography. The reasoning behind is that equality is in fact obtained asymptotically as $n \rightarrow \infty$, cf. (17) and Theorem 2. Alternatively, one can choose the ‘simplest’ function according to some complexity criterion, e.g. the TV-norm as in (15).

Global choice. In light of the last remark, our strategy is to determine the largest diffusivity – or the smoothest solution – such that the inequalities (22) hold. For the homogeneous diffusion process (7), this can be achieved algorithmically by starting with a large diffusivity $a \in \mathbf{R}^+$ which leads to oversmoothing, i.e. the corresponding estimator $\hat{f}_{\text{hom.diff.}}$ leads to residuals r_{ij} that violate the multiresolution criterion (22). Then, the diffusivity is reduced by a prespecified amount until the corresponding estimator gives residuals that fulfil the criterion. We refer to Section 5 for more details.

Davies and Kovac (2001) have shown that such a strategy leads to consistent estimators when combined with the one-dimensional taut string, i.e. when the global smoothing parameter is given by the number of extreme values of the taut string estimator $\hat{f} : \mathbf{R} \rightarrow \mathbf{R}$; an application to the selection of peaks in X-ray diffractograms is given in (Davies et al. 2008). Boysen et al. (2007, 2009) proved a similar consistency result for selecting the number of jumps in jump regression.

Dyadic squares partitioning. We yet have to specify what collection \mathcal{P}_n of subsets we will use; for the sake of brevity such collections will be called *partitionings*. Obviously, they do not constitute what is commonly referred to as a *partition* of a set, i.e. a disjoint composition of the latter, but they rather comprise several partitions at different scales as we shall see. Indeed, we want such a partitioning to consist of subsets that allow to detect deviations from the hypothesis at different resolutions, both at coarse and fine scales. However, it should not be chosen too rich either. From a statistical point of view, taking too many subsets results in too many multiple tests being performed which can be seen analytically from the asymptotics in (17) breaking down. Also, for each subset $P \in \mathcal{P}_n$ its multiresolution coefficient ω_P needs to be determined in our iterative procedure, rendering a large partitioning \mathcal{P}_n computationally very expensive. Furthermore it is

also necessary that the condition for each individual $P \in \mathcal{P}_n$ can be quickly checked. We now show how all these conditions can be met.

The subsequent theory is based on a *dyadic squares partitioning*, whereas for most practical purposes we use more subtle partitionings such as wedgelets or curvelets, see below. The dyadic squares partitioning originates from (Donoho 1997), cf. also (Kolaczyk et al. 2005) and (Antoniadis et al. 2009) for applications to image segmentation and classification, resp. It is obtained by splitting the image recursively into four equal subsquares until some pre-specified lowest scale is reached, cf. Figure 3 (left). This partitioning covers a wide range of scales with comparatively few subsets. At the same time it allows fast computation of the multiresolution coefficients through cumulative sums: define the matrix R of cumulative sums by

$$R = \left(\sum_{k=1}^i \sum_{l=1}^j r_{kl} \right)_{ij}. \quad (23)$$

From R , the multiresolution coefficient ω_P of a rectangle $P = [i_1, i_2] \times [j_1, j_2]$ can readily be obtained by

$$\begin{aligned} \sum_{(i,j) \in P} r(x_{ij}) &= R_{i_2, j_2} - R_{i_1-1, j_2} \mathbb{1}_{\{i_1 > 1\}} - R_{i_2, j_1-1} \mathbb{1}_{\{j_1 > 1\}} \\ &\quad + R_{i_1-1, j_1-1} \mathbb{1}_{\{i_1 > 1, j_1 > 1\}}. \end{aligned} \quad (24)$$

As promised in Section 3, we will now give the asymptotic distribution of M_n in (16) for a dyadic partitioning of a cube $K_n = \{1, \dots, n\}^d$ in arbitrary dimension d . First, let us recall the following well-known fact, see e.g. (Leadbetter et al. 1983, Theorem 1.5.3). If $\{\eta_i, i \in \mathbb{N}\}$ are independent standard Gaussian random variables, then for every $\tau \in \mathbf{R}$

$$\lim_{n \rightarrow \infty} \mathbf{P} \left[\max_{i=1, \dots, n} \eta_i \leq a_n + b_n \tau \right] = \exp(-e^{-\tau}), \quad (25)$$

where the normalising sequences a_n, b_n are defined in (21). As it turns out, we obtain the same asymptotic distribution, a *Gumbel* distribution in the case of a dyadic partitioning:

Theorem 2. *If \mathcal{P}_n is the collection of dyadic subcubes of K_n and $(r_i)_{i \in K_n}$ are independent standard Gaussian random variables, then for every $\tau \in \mathbf{R}$,*

$$\lim_{n \rightarrow \infty} \mathbf{P}[M_n \leq a_{\# \mathcal{P}_n} + b_{\# \mathcal{P}_n} \tau] = \exp(-e^{-\tau}), \quad (26)$$




Figure 3: Different scales of a dyadic squares partitioning (left), and with a line cutting through thereby creating wedgelets (right).

where

$$M_n = \max_{P \in \mathcal{P}_n} \frac{|\sum_{i \in P} r_i|}{\sqrt{\#P}}.$$

We note that other “dyadic-type” collections of scanning subsets (say, p -adic cubes, $p = 2, 3, \dots$) will also lead to this asymptotic distribution, as can be seen from the proof in the appendix.

Local smoothing. If we want to allow the diffusivity to vary locally as for the inhomogeneous diffusion introduced in Section 2, we have to modify the approach taken for the homogeneous diffusion above. It is the *locality* of the multiresolution criterion that allows for such a modification: if the criterion is violated, we can not only conclude that the diffusivity was too high but from (22) we can also infer *where*. Accordingly, each time a violation occurs for a multiresolution coefficient ω_P of some $P \in \mathcal{P}_n$, we reduce the local smoothing parameter $a(x)$ on that particular subset P only, keeping the diffusivity at its current value in the rest of the image.

As for the homogeneous diffusion, we start by initializing $a(x)$ to a large constant. Then we compute the estimator and its residuals, check the multiresolution criterion and adapt the smoothing parameter locally. This strategy has first been introduced by Meise (2004). Obviously, if the hypothesis is violated on a small subset P , we can *logically* conclude that the hypothesis must also be violated on any superset $P' \supset P$. Note that this is not to say

that if the *empirical criterion* is violated on a subset that it will also be violated on all supersets – this is clearly not true in general; however, if we decide against the hypothesis on such a subset we cannot accept the hypothesis on any superset. Hence we start by checking the multiresolution criterion on the smallest scale, considering larger sets only if no subset has shown a violation yet, thereby avoiding the diffusivity being reduced several times at the same spot. We then iterate the process until no further violations of the inequalities in (22) are detected, giving our final estimate.

Note that this results in the diffusivity being piecewise constant on the dyadic squares partitioning. In order to be able to adapt it to finer geometric features, we enhance the latter by adding so-called wedgelets. These are obtained by dividing a dyadic square into two parts by a straight line. We draw a certain set of lines through the image domain and add the resulting wedgelets to the partitioning. This is illustrated in Figure 3 (right). For a detailed description of *wedgelet partitionings* we refer to (Donoho 1999) and (Friedrich 2005). Fast computation of the multiresolution coefficients of wedgelets can be done by a similar though somewhat more involved argument as in (24); for details we refer to (Friedrich et al. 2007).

We use the wedgelet partitioning only to increase the flexibility when updating the diffusivity: if a violation is detected on a dyadic square P , all wedgelets subdividing that square are considered. If the criterion is fulfilled for all wedgelets contained in P the smoothing parameter will be reduced on P . Otherwise, the smoothing parameter will be reduced on the wedgelet W which yields the largest absolute value of the multiresolution coefficient ω_W . Again, all supersets of P will be ignored in that iteration. The final algorithm is described in Figure 4.

For this procedure to be implemented, we need an estimator of the noise's standard deviation σ . Based on the normality assumption, and taking the small number of pixels at sharp edges into account, we use a robust estimator based on the median of absolute differences, namely

$$\hat{\sigma} = \frac{1}{2\Phi^{-1}(0.75)} \text{Med} \{ |y_{i,j} - y_{i-1,j} - y_{i,j-1} + y_{i-1,j-1}| : i, j = 2, \dots, n \}, \quad (28)$$

where Φ denotes the cumulative distribution function of $\mathcal{N}(0, 1)$, cf. (Davies and Kovac 2001). For smooth signals, polynomially weighted, difference-based estimators might be more appropriate, see (Munk et al. 2005); note that our estimator is indeed unbiased if f is affine.

- choose some $\delta > 0$, e.g. such that the error of the first kind is 5% which can be obtained from simulations
- determine σ
- initialize $a : \{1, \dots, n\}^2 \rightarrow \mathbf{R}^+$ to a large constant

loop

- obtain current reconstruction $\hat{f}(\cdot) = L_a y$, see (8)
- compute residuals $r_{ij} = y_{ij} - \hat{f}_{ij}$
- for each dyadic square P determine its multiresolution coefficient

$$\omega_P = \frac{1}{\sqrt{\#\{x_{ij} \in P\}}} \sum_{x_{ij} \in P} r_{ij} \quad (27)$$

if there is some P on which the criterion is violated,
i.e. having $|\omega_P| > \sigma \sqrt{\delta \log n^2}$, **then**

for all such P **do**

if there is some subset $P' \subset P$ with a violation **then**

- do nothing

else

- determine the multiresolution coefficients ω_W of all wedgelets W comprising this square

if $\max_W |\omega_W| > \sigma \sqrt{\delta \log n^2}$, i.e. there is some wedgelet with a large coefficient, **then**

- reduce a on the wedgelet $W = \operatorname{argmax}_W \omega_W$

else

- reduce a on P

end if

end if

end for

else

- stop and return the current estimate \hat{f} , i.e. the first estimate without a violation

end if

end loop

Figure 4: Algorithm to determine local diffusivity and hence the reconstructed image.

The algorithms described in this section clearly do not depend on the kind of penalty used, cf. (9). All that is needed is a global or local smoothness parameter a , corresponding to the homogeneous or inhomogeneous diffusivity. In particular, we proceed in the same way for the TV regularization in (10). See (Davies and Meise 2008) for an application of this approach to one- and two-dimensional weighted smoothing splines.

5 Numerical details and results

In this section we show simulations where the noiseless image f of (1) is the 256×256 pixel test image shown in Figure 1(a), whose values spread over $[0, 5]$. To this we added Gaussian noise according to (2) with $\sigma = 1$, resulting in a signal-to-noise ratio of 5, see Figure 1(b).

In each iteration, the smoothing parameter was reduced where necessary through multiplication with a factor $\lambda < 1$ as described in Section 4. In order to speed up the algorithm we chose λ according to the size of the multiresolution coefficient: the larger ω_P , the smaller λ was chosen, taking care not too quickly to reduce the diffusivity which would result in undersmoothing. Note that the inhomogeneous diffusion estimator can be found efficiently by solving (8) since the discretized Laplacian is given by a sparse band matrix, rendering Gauss-Seidel iterations an appropriate solution method.

The algorithms were implemented in Matlab in a modular fashion, thus allowing to use any localized regression method as discussed in Section 2. Figure 5 shows results for homogeneous and inhomogeneous diffusion. Clearly, in order to reconstruct fine details and sharp edges well, such that the multiresolution conditions (22) are satisfied, a very small diffusivity ($a = 0.934$) is needed when chosen globally; this results in considerable undersmoothing elsewhere in Figure 5(a). Choosing the diffusivity locally resolves this difficulty, allowing e.g. to strongly smooth large areas of the background or the circle in the upper left while not compromising small level details like the dots in the lower right or any sharp edges, see Figures 5(b) and 8(a). This behaviour of the local diffusivity as an edge-detector is very much apparent in Figure 7(a).

For comparison, we also show results for the total-variation (TV) regularization introduced in (10), again both for global and local smoothing parameters; see Figure 6(a) and (b), respectively. The difficulty when choosing a global smoothing parameter is again well visible but the local smoothing




Figure 5: Results of homogeneous (a) and inhomogeneous (b) diffusion.




Figure 6: Results for TV regularization with global (a) and local (b) choice of the smoothing parameter.

parameter appears to have more difficulty adapting to the test object, permitting considerable undersmoothing on the larger dots in the lower right, cf. Figures 7(b) and 8(b). The TV penalty was implemented by approximation with a differentiable functional, as described e.g. by Vogel (2002).

When the signal-to-noise ratio is reduced to 2, see the simulated data in Figure 9(a), the reconstruction gets smoothed more strongly (b), since the residuals no longer carry enough statistically significant information to allow the oversmoothing of the smaller features to be detected, though still the




Figure 7: Local smoothing parameters for diffusion (a) and TV penalty (b).

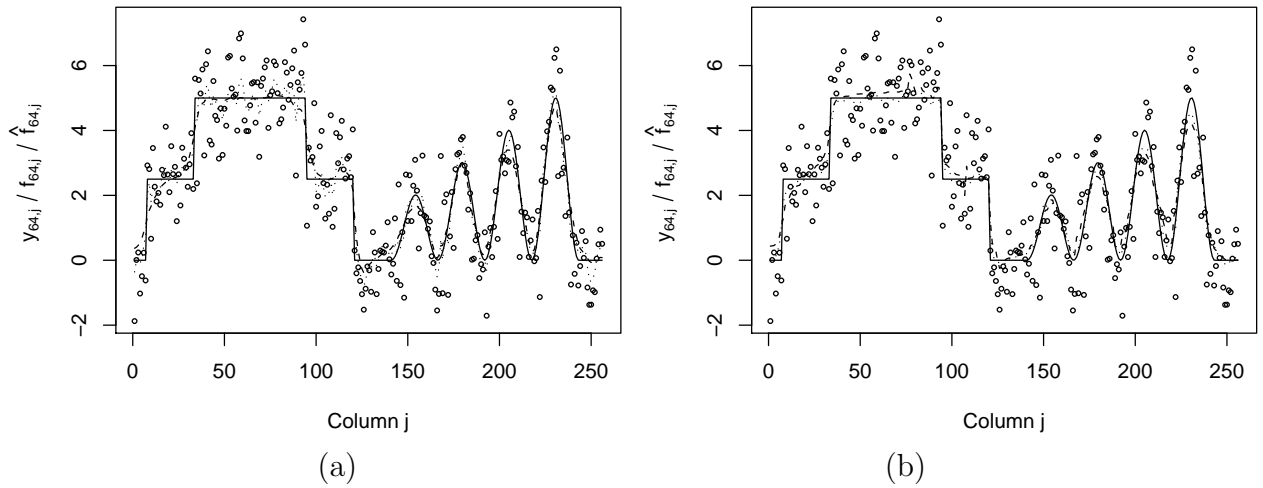


Figure 8: Cut along row 64 indicated in Figure 1(a) for diffusion (a) and TV penalty (b); solid lines gives the true f , points the data y , dotted lines correspond to a global and dashed lines to a local choice of the smoothing parameter.

only feature not distinguishable from the noise is the smallest of the nine circles in the lower right.

Poisson data with high intensities can also be treated with this methodology as remarked at the end of Section 3; instead of determining a constant variance at the beginning of the algorithm, one uses the local variance pre-




Figure 9: (a) Simulated data y with noise level $\sigma = 2.5$. (b) Corresponding reconstruction using inhomogeneous diffusion.




Figure 10: (a) Simulated Poisson data y with intensities in $[50, 100]$. (b) Corresponding reconstruction using inhomogeneous diffusion.

dicted by the reconstruction to normalize the residuals. We also simulated this situation for intensities within $[50, 100]$, i.e. again with a signal-to-noise ratio of 5, see Figure 10. The reconstruction demonstrates the applicability of our approach also in this situation.

We conclude this section with an application where the image has been obtained by a CCD camera attached to a confocal microscope, see Figure 11, data courtesy of Emre Togan, Department of Applied Physics, Harvard Uni-




Figure 11: (a) Poisson data of photoluminescence. (b) Corresponding reconstruction using inhomogeneous diffusion.

versity. The data set shows photoluminescence in a diamond sample, where high photoluminescence marks the so-called nitrogene-vacancy centres in the diamond. These are of major interest to quantum information science where they have been proposed as qubits, i.e. for storage, as they form a solid-state system whose spins can be manipulated at room temperature. This application therefore aims at removing noise from the image in order to aid the researcher in detecting these nitrogene-vacancy centres, such that subsequent experiments can be conducted on them, see (Dutt et al. 2007) and the references therein. We note that our reconstruction reduces the noise considerably while keeping small-scale features, much to the satisfaction of the physicists involved.

6 Discussion

We have demonstrated that the statistical multiresolution criterion allows to determine both global and local smoothing parameters in a fully data-driven procedure. For inhomogeneous diffusion, it acts as an edge-detector – quite as expected: at edges, only very little smoothing is to be permitted for the criterion to be fulfilled there, while large regions that are well approximated by their first order Taylor series expansion may be smoothed strongly.

It appears that using a penalty that is more capable of dealing with

sharp edges, like TV regularization, does not improve upon the inhomogeneous diffusion, or even performs worse. A possible explanation is that the multiresolution criterion itself is able to detect edges, thus eliminating the major drawback of the diffusion process – whilst keeping its advantageous properties in smooth areas. TV regularization with a global parameter, however, already is capable of reconstructing sharp edges and can localize them, thus it cannot benefit as much from the multiresolution criterion’s ability to choose the parameter locally, especially as the smoothness of the reconstruction cannot be improved.

We emphasize once again that the multiresolution analysis can be extended to other error distributions as well, see Section 3. Also, generalizations to higher dimensions are possible; in particular, three-dimensional images or movies could be treated efficiently, too.

Acknowledgements

T. Hotz and P. Marnitz gratefully acknowledge support by the German Federal Ministry of Education and Research, Grant 03MUPAH6. A. Munk and Z. Kabluchko acknowledge support by the German Research Foundation’s FOR 916, and A. Munk also by SFB 755.

A Proofs

The proof of Theorem 1 on page 12 will be based on the following lemma which is known as Berman’s Inequality, see e.g. (Leadbetter et al. 1983, Theorem 4.2.1).

Lemma 1. *Let (ξ_1, \dots, ξ_N) be a Gaussian vector with standard margins and covariance matrix $(\rho_{ij})_{i,j=1,\dots,N}$, and let η_1, \dots, η_N be independent standard Gaussian variables. Then, for every $u > 0$,*

$$\mathbf{P} \left[\max_{1 \leq i \leq N} \xi_i \leq u \right] - \mathbf{P} \left[\max_{1 \leq i \leq N} \eta_i \leq u \right] \leq \frac{1}{2\pi} \sum_{1 \leq i < j \leq N} \frac{|\rho_{ij}|}{\sqrt{1 - \rho_{ij}^2}} \exp \left(-\frac{u^2}{1 + \rho_{ij}} \right).$$

Proof of Theorem 1 on page 12. Let $u_N = a_N + b_N \tau$, with a_N and b_N

as in (21). We then have, using Berman's Inequality,

$$\begin{aligned} & \mathbf{P} \left[\max_{1 \leq i \leq N} \xi_i^{(N)} \leq u_N \right] - \mathbf{P} \left[\max_{1 \leq i \leq N} \eta_i \leq u_N \right] \\ & \leq O(1) \sum_{1 \leq i < j \leq N} |\text{Cov}(\xi_i^{(N)}, \xi_j^{(N)})| \exp \left(-\frac{u_N^2}{1 + \rho} \right) \\ & \leq O(N^{1+\varepsilon}) \exp \left(-\frac{u_N^2}{1 + \rho} \right) \end{aligned}$$

as $N \rightarrow \infty$. Noting that $u_N \sim \sqrt{2 \log N}$ and choosing ε small enough, we see that the right-hand side is $o(1)$ as $N \rightarrow \infty$. Combining this estimate with (25), we obtain (20), which finishes the proof of the theorem. \square

Proof of Theorem 2 on page 15. We are going to apply Theorem 1 with $N = \#\mathcal{P}_n$ to the random vector $(\xi_P^{(N)})_{P \in \mathcal{P}_n}$, where

$$\xi_P^{(N)} = \frac{\sum_{i \in P} r_i}{\sqrt{\#\mathcal{P}}}, \quad P \in \mathcal{P}_n.$$

First we show that (19) holds. Let P_1 and P_2 be two different dyadic subcubes of K_n . If $P_1 \cap P_2 = \emptyset$, then we have $\text{Cov}(\xi_{P_1}^{(N)}, \xi_{P_2}^{(N)}) = 0$, so that (19) is fulfilled. If $P_1 \cap P_2 \neq \emptyset$, then we have either $P_1 \subset P_2$ or $P_2 \subset P_1$. Assuming that, say, $P_1 \subset P_2$, we obtain (due to the dyadic structure) that $\#P_1 \leq \#P_2/2$, and hence,

$$\text{Cov}(\xi_{P_1}^{(N)}, \xi_{P_2}^{(N)}) = \frac{\#(P_1 \cap P_2)}{\sqrt{\#P_1 \#P_2}} = \frac{\sqrt{\#P_1}}{\sqrt{\#P_2}} \leq \frac{1}{\sqrt{2}}.$$

Now we show that (18) holds. A dyadic cube with side length 2^k contains $2^{d(k-i)}$ dyadic cubes with side length 2^i , where $0 \leq i \leq k$. Therefore, the total number of dyadic cubes which are contained in a dyadic cube with side length 2^k is $\sum_{i=0}^{k-1} 2^{d(k-i)} \leq 2^{d(k+1)}$. Further, the number of dyadic cubes in \mathcal{P}_n having side length 2^k is $\lfloor n/2^k \rfloor^d$. Note also that, trivially, $\#\mathcal{P}_n \geq n^d$. So, we may estimate the cardinality on the left-hand side of (18) from above by

$$2 \sum_{k=0}^{\lfloor \log n / \log 2 \rfloor} \lfloor n/2^k \rfloor^d \cdot 2^{d(k+1)} = O(n^d \log n) = O(\#\mathcal{P}_n^{1+\varepsilon})$$

as $N \rightarrow \infty$ and for every $\varepsilon > 0$. This finishes the proof. \square

References

- Alvarez, L., Guichard, F., Lions, P.-L., and Morel, J.-M. (1993). Axioms and fundamental equations of image processing. *Arch. Rational Mech. Anal.*, 123:199–257.
- Antoniadis, A., Bigot, J., and von Sachs, R. (2009). A multiscale approach for statistical characterization of functional images. *Journal of Computational and Graphical Statistics*, 18:216–237.
- Aubert, G. and Kornprobst, P. (2002). *Mathematical problems in image processing: partial differential equations and the calculus of variations*. Number 147 in Applied mathematical sciences. Springer, New York.
- Besag, J. (1986). On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(2):259–302.
- Bissantz, N., Hohage, T., Ruymgaart, F., and Munk, A. (2007). Convergence rates of general regularization methods for statistical inverse problems and applications. *SIAM J. Numerical Analysis*, 45:2610–2636.
- Bissantz, N., Mair, B., and Munk, A. (2006). A multi-scale stopping criterion for MLEM reconstructions in PET. *IEEE Nucl. Sci. Symp. Conf. Rec.*, 6:3376–3379.
- Bissantz, N., Mair, B., and Munk, A. (2008). A statistical stopping rule for MLEM reconstructions in PET. *IEEE Nucl. Sci. Symp. Conf. Rec.*, 8:4198–4200.
- Boysen, L., Kempe, A., Munk, A., Liebscher, V., and Wittich, O. (2009). Consistencies and rates of convergence of jump penalized least squares estimators. *Annals of Statistics*, 37:157–183.
- Boysen, L., Liebscher, V., Munk, A., and Wittich, O. (2007). Scale space consistency of piecewise constant least squares estimators – another look at the regressogram. In Cator, E. A., Jongbloed, G., Kraaikamp, C., Lopuhaä, H. P., and Wellner, J. A., editors, *Asymptotics: Particles, Processes and Inverse Problems: Festschrift for Piet Groeneboom*, volume 55 of *IMS Lecture Notes*, pages 65–84. Institute of Mathematical Statistics, Beachwood, Ohio.

- Buades, A., Coll, B., and Morel, J. (2005). A review of image denoising algorithms, with a new one. *Multiscale Modeling and Simulation (SIAM interdisciplinary journal)*, 4(2):490–530.
- Chaudhuri, P. and Marron, J. S. (2000). Scale space view of curve estimation. *The Annals of Statistics*, 28(2):408–428.
- Chu, C. K., Glad, I. K., Godtlielsen, F., and Marron, J. S. (1998). Edge-preserving smoothers for image processing. *Journal of the American Statistical Association*, 93(442):526–541.
- Davies, P. L., Gather, U., Meise, M., Mergel, D., and Mildenerger, T. (2008). Residual-based localization and quantification of peaks in X-ray diffractograms. *Annals of Applied Statistics*, 2(3):861–886.
- Davies, P. L. and Kovac, A. (2001). Local extremes, runs, strings and multiresolution. *The Annals of Statistics*, 29(1):1–65.
- Davies, P. L., Kovac, A., and Meise, M. (2009). Nonparametric regression, confidence regions and regularization. *Annals of Statistics*. To appear.
- Davies, P. L. and Meise, M. (2008). Approximating data with weighted smoothing splines. *Journal of Nonparametric Statistics*, 20(3):207 – 228.
- Devroye, L. and Lugosi, G. (2001). *Combinatorial Methods in Density Estimation*. Springer Verlag, New York.
- Donoho, D. L. (1997). CART and best-ortho-basis: a connection. *The Annals of Statistics*, 25(5):1870–1911.
- Donoho, D. L. (1999). Wedgelets: Nearly-minimax estimation of edges. *The Annals of Statistics*, 27(3):859–897.
- Dümbgen, L. and Spokoiny, V. (2001). Multiscale testing of qualitative hypotheses. *The Annals of Statistics*, 29(1):124–152.
- Dümbgen, L. and Walther, G. (2008). Multiscale inference about a density. *The Annals of Statistics*, 36(4):1758–1785.
- Dutt, M. V. G., Childress, L., Jiang, L., Togan, E., Maze, J., Jelezko, F., Zibrov, A. S., Hemmer, P. R., and Lukin, M. D. (2007). Quantum Register Based on Individual Electronic and Nuclear Spin Qubits in Diamond. *Science*, 316(5829):1312–1316.

- Ern, A. and Guermond, J.-L. (2004). *Theory and practice of finite elements*, volume 159 of *Applied mathematical sciences*. Springer, New York.
- Friedrich, F. (2005). *Complexity penalized segmentations in 2D*. PhD thesis, Technical University of Munich.
- Friedrich, F., Demaret, L., Führ, H., and Wicker, K. (2007). Efficient moment computation over polygonal domains with an application to rapid wedgelet approximation. *SIAM J. Scientific Computing*, 29(2):842–863.
- Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741.
- Genovese, C. and Wasserman, L. (2008). Adaptive confidence bands. *The Annals of Statistics*, 36(2):875–905.
- Gijbels, I. and Mammen, E. (1998). Local adaptivity of kernel estimates with plug-in local bandwidth selectors. *Scandinavian Journal of Statistics*, 25:503–520.
- Guichard, F., Morel, J.-M., and Ryan, R. (2004). Contrast invariant image analysis and PDE's. Unpublished book available at <http://www.cmla.ens-cachan.fr/fileadmin/Membres/morel/JMMBookOct04.pdf>.
- Hall, P. and Titterton, D. M. (1986). On some smoothing techniques used in image restoration. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(3):330–343.
- Kabluchko, Z. (2007). Extreme-value analysis of standardized increments. Preprint available at <http://www.arxiv.org/abs/0706.1849>.
- Kabluchko, Z. and Munk, A. (2008). Shao's theorem on the maximum of standardized random walk increments for multidimensional arrays. *ESAIM Prob. Stat.* to appear.
- Koenderink, J. J. (1984). The structure of images. *Biological Cybernetics*, 50:363–370.
- Kolaczyk, E. D., Ju, J., and Gopal, S. (2005). Multiscale, multigranular statistical image segmentation. *Journal of the American Statistical Association*, 100(472):1358–1369.

- Korostelev, A. and Tsybakov, A. (1993). *Minimax Theory of Image Reconstruction*, volume 82 of *Lecture Notes in Statistics*. Springer, New York.
- Leadbetter, M., Lindgren, G., and Rootzén, H. (1983). *Extremes and related properties of random sequences and processes*. Springer Series in Statistics. Springer-Verlag, New York.
- Lepskii, O. V. (1990). On a problem of adaptive estimation in Gaussian white noise. *Theory of Probability and its Applications*, 35(3):454–466.
- Lepskii, O. V., Mammen, E., and Spokoiny, V. G. (1997). Optimal spatial adaptation to inhomogeneous smoothness: an approach based on kernel estimates with variable bandwidth selectors. *The Annals of Statistics*, 25(3):929–947.
- Li, K.-C. (1989). Honest confidence regions for nonparametric regression. *The Annals of Statistics*, 17(3):1001–1008.
- Mammen, E. and van de Geer, S. (1997). Local adaptive regression splines. *The Annals of Statistics*, 25:387–413.
- Meise, M. (2004). *Residual Based Selection of Smoothing Parameters*. PhD thesis, Universität Duisburg-Essen.
- Mrázek, P., Weickert, J., and Bruhn, A. (2006). On robust estimation and smoothing with spatial and tonal kernels. In Klette, R., Kozera, R., Noakes, L., and Weickert, J., editors, *Geometric Properties for Incomplete data*, volume 31 of *Computational Imaging and Vision*, pages 335–352. Springer Netherlands.
- Munk, A., Bissantz, N., Wagner, T., and Freitag, G. (2005). On difference-based variance estimation in nonparametric regression when the covariate is high dimensional. *Journal of the Royal Statistical Society, Series B*, 67(1):19–41.
- Nardi, Y., Siegmund, D. O., and Yakir, B. (2008). The distribution of maxima of approximately Gaussian random fields. *The Annals of Statistics*, 36(3):1375–1403.
- Pollard, D. (1984). *Convergence of Stochastic Processes*. Springer-Verlag, New York.

- Polzehl, J. and Spokoiny, V. (2000). Adaptive weights smoothing with applications to image restoration. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 62(2):335–354.
- Polzehl, J. and Spokoiny, V. (2003). Image denoising: Pointwise adaptive approach. *The Annals of Statistics*, 31(1):30–57.
- Qiu, P. (2005). *Image processing and jump regression analysis*. Wiley series in probability and statistics. Wiley-Interscience, Hoboken, NJ.
- Qiu, P. (2007). Jump surface estimation, edge detection and image restoration. *Journal of the American Statistical Association*, 102(478):745–756.
- Rudin, L. I., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1–4):259–268.
- Scherzer, O., Grasmair, M., Grossauer, H., Haltmeier, M., and Lenzen, F. (2009). *Variational Methods in Imaging*, volume 167 of *Applied Mathematical Sciences*. Springer, Berlin.
- Shao, Q.-M. (1995). On a conjecture of Révész. *Proceedings of the American Mathematical Society*, 123(2):575–582.
- Siegmund, D. O. and Venkatraman, E. S. (1995). Using the generalized likelihood ratio statistic for sequential detection of a change-point. *The Annals of Statistics*, 23(1):255–271.
- Siegmund, D. O. and Worsley, K. J. (1995). Testing for a signal with unknown location and scale in a stationary Gaussian random field. *The Annals of Statistics*, 23(2):608–639.
- Siegmund, D. O. and Yakir, B. (2000). Tail probabilities for the null distribution of scanning statistics. *Bernoulli*, 6(2):191–213.
- Stichtenoth, R. (2007). *Signal and Image Denoising Using Inhomogenous Diffusion*. PhD thesis, Universität Duisburg-Essen, Essen.
- Tsybakov, A. B. (1998). Pointwise and sup-norm sharp adaptive estimation of functions on the Sobolev classes. *The Annals of Statistics*, 26(6):2420–2469.

- Vogel, C. R. (2002). *Computational methods for inverse problems*, volume 23 of *Frontiers in applied mathematics*. SIAM, Philadelphia.
- Wand, M. P. and Jones, M. C. (1995). *Kernel Smoothing*, volume 60 of *Monographs on statistics and applied probability*. Chapman & Hall, London.
- Weickert, J. (1998). *Anisotropic Diffusion in Image Processing*. European Consortium for Mathematics in Industry. B. G. Teubner, Stuttgart.
- Winkler, G. (2003). *Image analysis, random fields and Markov chain Monte Carlo methods*. Number 27 in *Applications of mathematics*. Springer, 2nd edition.
- Witkin, A. P. (1983). Scale-space filtering. In *Proc. 8th Int. Joint Conf. Art. Intell.*, pages 1019–1022, Karlsruhe, Germany.